



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 714 063 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention

of the grant of the patent:

20.02.2002 Bulletin 2002/08

(51) Int Cl.7: G06F 11/32, G06F 11/34

(21) Application number: 95111110.3

(22) Date of filing: 14.07.1995

(54) Method for identifying ways to improve performance in computer data storage systems

Verfahren zum Aufzeigen von Wegen zur Verbesserung der Leistung von Computerdatenspeichersystemen

Procédé pour identifier des chemins pour améliorer la performance dans des systèmes de mémorisation de données d'ordinateur

(84) Designated Contracting States:

DE FR GB

(30) Priority: 22.11.1994 US 343809

(43) Date of publication of application:

29.05.1996 Bulletin 1996/22

(73) Proprietor: Hewlett-Packard Company, A
Delaware Corporation
Palo Alto, CA 94304 (US)

(72) Inventors:

- Voigt, Douglas L.
Boise, Idaho 83702 (US)
- Nelson, Marvin D.
Boise, Idaho 83704 (US)
- Dolphin, Wade A.
Boise, Idaho 83704 (US)

(74) Representative: Schoppe, Fritz, Dipl.-Ing.

Schoppe, Zimmermann & Stöckeler

Patentanwälte Postfach 71 08 67

81458 München (DE)

(56) References cited:

EP-A- 0 259 224

EP-A- 0 518 574

WO-A-94/09429

- PROCEEDINGS OF THE CONFERENCE ON ARTIFICIAL INTELLIGENCE APPLICATIONS, MIAMI BEACH, FEB. 24 - 28, 1991, vol. VOL. 1, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, XP 000299010 SE YOUNG PARK ET AL 'OASIX: A REAL-TIME KNOWLEDGE-BASED SYSTEM FOR UNIX OPERATIONS AND ADMINISTRATION'
- IBM TECHNICAL DISCLOSURE BULLETIN, vol. 35, no. 2, 1 July 1992 pages 92-94, XP 000313231 'USER INTERFACE FOR PERFORMANCE MONITOR'

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 0 714 063 B1

Description**FIELD OF THE INVENTION**

5 [0001] This invention relates to data storage systems, such as disk array data storage systems, and methods for identifying techniques for improving their performance based upon demands imposed by a user.

BACKGROUND OF THE INVENTION

10 [0002] The performance of data storage systems inevitably varies over the course of a lifetime. Performance is affected by several different factors, including the types of components employed in the storage system, the operational settings, and a customer's storage requirements, such as capacity and I/O demand. As components fail or become over- or under-utilized, performance might suffer. If the operational settings are improper for the customer's storage configuration, performance might be degraded. If the customer imposes storage demands which are too rigorous for 15 the present system, resulting in overburdening the storage system, performance will be adversely affected.

[0003] These performance-affecting factors vary during the life of the data storage system. Moreover, the factors might vary dramatically in the short term. When performance noticeably changes, the user is often tempted to make 20 system alterations that are expected or believed to improve performance. However, the alterations might be based on perceived problems which are actually symptoms of a more fundamental hidden problem. This type of uninformed problem solving might lead the user to incur unnecessary costs and system down time without ever satisfactorily improving performance. For instance, the user might believe that adding another disk drive to the disk array would improve 25 performance, when in reality, performance could be improved by simply reconfiguring the data storage system in a new manner which better accommodates the user's evolving storage demands.

[0004] It would therefore be advantageous to evaluate system performance during operation in an effort to monitor 25 or anticipate situations which adversely impact performance. Once detected, the customer could then be alerted that the system is not performing optimally, and instructed as to what could be done to improve performance.

[0005] EP-A-0518574 relates to a graphical system resource monitor for depicting, in real time, a data processing 30 system's internal system resource utilization, such as RAM utilization. During operation of the system resource monitor, key performance data for various resources are tracked. A graphical interface is provided in order to depict both instantaneous and recent/past records of resource utilization.

[0006] EP-A-0259224 relates to a system for performing a data processor performance evaluation including data collection, analysis and report generation. The data collection involves the periodic selection of measurable values 35 (metrics) and system parameters, the analysis involves the application of certain rules comparing metrics and parameters to corresponding thresholds and the report generation includes the display of recommendations to be taken when certain rules have triggered a certain number of times.

[0007] It is the object of the present invention to provide methods for flexibly identifying areas for performance improvement in a data storage system.

[0008] This object is achieved by methods according to claims 1 and 5.

[0009] According to one aspect of this invention, a method is provided for identifying areas for improvement in a data 40 storage system. Multiple performance metrics are defined which can be measured during operation of the data storage system. Example performance metrics include the total number of blocks read by the host, the total number of blocks written by the host, and the amount of time during which commands were outstanding from the host. A specification for the expected values of performance metrics is also provided. When the performance metrics are in compliance with the specification, the data storage system is operating at optimal performance.

[0010] According to this method, the performance metrics are sampled during operation of the data storage system. During operation, a performance history is constructed to indicate how the data storage system is performing over a period of time. The performance history examines the activities and storage demands imposed by the user, such as number of I/Os per second, memory transactions, used and unused memory space, etc. The performance history is presented to the user, preferably in the form of a graph which plots the activities and storage demands over time.

[0011] When inspecting the performance history, the user can select an evaluation interval for closer scrutiny. This is preferably accomplished by graphically overlaying an interval region on the performance history and then permitting the user to select the evaluation interval by graphically manipulating the interval region to overlay the desired time frame. For the selected evaluation interval, it can be determined whether or not any performance metric is in compliance with the specification. In the event that one or more performance metrics are not in compliance with the specification, 50 one or more various suggestions are provided to help improve performance of the data storage system. The suggestions are chosen based upon which of the performance metrics are not in compliance with the specification.

[0012] Through the use of these suggestions, the user can make informed decisions on how to improve performance of the data storage system. The performance planning method of this invention also helps the user anticipate future

storage requirements and timely upgrade the data storage system to accommodate them.

DESCRIPTION OF THE DRAWINGS

5 [0013] Preferred embodiments of the invention are described below with reference to the following accompanying drawings depicting examples embodying the best mode for practicing the invention.

[0014] Fig. 1 is a diagrammatical illustration of a host computer station connected to a data storage system of this invention.

[0015] Fig. 2 is a block diagram of a data storage system embodied as a hierarchical disk array.

10 [0016] Fig. 3 is a flow diagram depicting a method for identifying areas for performance improvement in the data storage system.

[0017] Fig. 4 is a model illustrating data flow through the data storage system, and locations where performance metrics can be measured.

[0018] Fig. 5 is a visual screen display showing a performance history of the data storage system.

15 [0019] Fig. 6 is a visual screen display showing performance planning suggestions.

[0020] Fig. 7 is a visual screen display showing the measured performance metrics.

DETAILED DESCRIPTION OF THE INVENTION

20 [0021] The methods of this invention are described in the context of computer disk array data storage systems which employ multiple disk drives. However, the principles of this invention are applicable to other types of data storage systems, such as reel-to-reel storage systems, tape back-up systems, and CD storage systems.

[0022] Fig. 1 shows a computer system 10 having a host computer terminal or station 12 connected to a data storage system 14 via I/O interface bus 16. Host computer station 12 includes a visual display monitor 18, a central processing unit (CPU) 20, and a keyboard 22.

25 [0023] Fig. 2 shows an example construction of data storage system 14 embodied as a hierachic disk array data storage system. Storage system includes disk array 30 having a plurality of storage disks 32, a disk array controller 34 coupled to the disk array 30 to coordinate data transfer to and from the storage disks 32, and a RAID management system 36.

30 [0024] The term "RAID" (Redundant Arrays of inexpensive Disks) means a disk array in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The redundant information enables regeneration of user data in the event that one of the array's member disks or the access path to it fails. A more detailed discussion of RAID systems is found in a book entitled, *The RAIDBook: A Source Book for RAID Technology*, published June 9, 1993, by the RAID Advisory Board, Lino Lakes, Minnesota.

35 [0025] Disk array controller 34 is coupled to disk array 30 via one or more interface buses 38, such as a small computer system interface (SCSI). RAID management system 36 is operatively coupled to disk array controller 34 via an interface protocol 40. The controller 34 is coupled to the host computer via I/O interface bus 16. RAID management system 36 can be embodied as a separate component, or configured within disk array controller 34 or within the host computer to provide a data manager means for controlling disk storage and reliability levels, and for transferring data among various reliability storage levels. These reliability storage levels are preferably mirror or parity redundancy levels, but can also include a reliability storage level with no redundancy at all.

40 [0026] The disk array controller 34 is implemented as a dual controller consisting of disk array controller A 34a and disk array controller B 34b. Dual controllers 34a and 34b enhance reliability by providing continuous backup and redundancy in the event that one controller becomes inoperable.

45 [0027] Storage disks 32 have example sizes of one to three Gigabytes. The storage disks can be independently connected or disconnected to mechanical bays that provide interfacing with SCSI bus 38. In one implementation, the data storage system is designed with twelve active mechanical bays. Four SCSI buses are used to interface these bays with disk array controller 34 (i.e., one bus per 3 mechanical bays). If the active bays are fully loaded, the data storage system has an example combined capacity of 12-36 Gigabytes.

50 [0028] The hierachic disk array 30 can be characterizable as different storage spaces, including its physical storage space and one or more virtual storage spaces. These various views of storage are related through mapping techniques. For example, the physical storage space of the disk array can be mapped into a RAID-level virtual storage space which delineates storage areas according to the various data reliability levels. For instance, some areas within the RAID-level virtual storage space can be allocated for a first reliability storage level, such as mirror or RAID level 1, and other areas can be allocated for a second reliability storage level, such as parity or RAID level 5. The RAID-level virtual view can be mapped to a second application-level virtual storage space which presents a contiguously addressable storage space. The physical configuration and RAID view of the storage space are hidden from the application view, which is presented to the user.

[0029] A memory map store 42 provides for persistent storage of the virtual mapping information used to map different storage spaces into one another. The memory mapping information can be continually or periodically updated by the controller or RAID management system as the various mapping configurations among the different views change. The memory map store 42 is embodied as two non-volatile RAMs (Random Access Memory) 42a and 42b, such as battery-backed RAMs, which are located in respective controllers 34a and 34b. The dual NVRAMs 42a and 42b provide for redundant storage of the memory mapping information.

[0030] The storage disks provide a physical storage capacity which is available to the host computer for data storage. It is desirable that this capacity be managed effectively to optimize performance, both in terms of efficiency of operation as well as satisfying the storage demands imposed by the user. Performance management techniques of this invention 10 are grounded in finding answers to difficult, imprecise questions which continuously evaluate how well the data storage system is performing. The following seven inquiries demonstrate an example set of performance decision making goals:

1. Is the data storage system a performance bottleneck?
2. Do current system configuration settings enable optimal performance? Example settings include SCSI transfer 15 rate and size of the logical units.
3. Would rebalancing storage disks on four SCSI channels 38 improve performance?
4. Would larger, faster, or additional storage disks significantly improve performance?
5. Would additional RAM in the disk array controller 34 significantly improve performance?
6. Would additional data storage systems 14 significantly improve performance?
7. Would upgrading to a new generation of disk array controller 34 or storage disk 32 improve performance?

[0031] The first enumerated decision making inquiry indicates that something is affecting the system, although the source may be indeterminable. Inquiries 2-5 involve making changes to the existing data storage system. Inquiries 6-7 25 involve alterations which involve additional components, such as modifications made after the capacity of one data storage system is fully consumed.

[0032] Performance improvement questions are difficult to answer with absolute accuracy as they involve some degree of prediction. However, a complete quantitative estimation of performance parameters requires high storage and analytic processing costs. The methods of this invention therefore focus on utilization of the existing components in the data storage system. Utilization data helps isolate performance problems, and helps identify components that 30 are operating well or poorly in their present configuration. The basic performance improvement strategy, then, is to invest in changes that will solve current problems, and not in changes to components that are already operating within specification.

[0033] The methods of this invention identify potential areas for performance improvement. The methods inform the user as to substandard performance, and generate suggestions which assist the user in rectifying the problems to 35 return the system to optimal performance.

[0034] Fig. 3 shows the preferred steps of a method for identifying areas for performance improvement. The initial step 50 is to define one or more performance metrics that can be measured during operation of the data storage system. Example performance metrics include the total number of blocks read by the host, the total number of blocks written 40 by the host, the amount of time during which commands were outstanding from the host, and the maximum SCSI wait queue size per device. A specification which establishes the standards for the performance metrics is also provided (step 52). The data storage system is operating at optimal performance when the performance metrics are in compliance with the specification, or said another way, when the performance metrics are "in spec". The standards established by the specification may be in the form of acceptance or tolerance computations with thresholds for various performance metrics.

[0035] Once the performance metrics and corresponding specification are defined, the data storage system is 45 equipped with means for monitoring the performance metrics. Such means is preferably implemented as firmware in the disk array controller 34 to extract component utilization data during operation.

[0036] Once a performance metric measurement means is installed, the next step 54 is to sample the performance metrics during operation of the data storage system. This sampling occurs in either the data storage system 14 or the host computer station 12. Fig. 4 shows an example model 80 of the controller firmware demonstrating the points at 50 which performance metrics can be sampled. The circles represent stages of I/O processing and the three-sided boxes represent queues. The model shows information flow from the host computer through to the storage disks. Requests made by the host flow first into a front end SCSI 82, and are then temporarily held in host queue 84. Decoder 86 dissects the requests into smaller commands for further processing. The commands are placed in cache queue 88 55 until cache 90 is ready to accept them. Workload is transformed in hierachic manager 92 in response to the RAID policies 94 regarding hierarchy and redundancy management. Flow continues to SCSI queue 96 which temporarily holds items for back end SCSI 98.

[0037] Performance metrics are extracted at various points along the model. For instance, the performance metric

for the total number of blocks read by the host is extracted at point A in front end SCSI 82. The following table 1 lists the example performance metrics, the sample points for the performance metrics, and the decision affected by the performance metrics.

5

TABLE 1

Metric	Sample Point	Decision Supported
Active - Amount of time during which >0 commands were outstanding from host	A	1) Is there a Bottleneck? 7) Upgrade improve performance?
SCSI Queue Max - Maximum SCSI wait queue size per device	D	3) Rebalance Disks? 4) Additional Disks Help?
Host Read Blocks - Total number of blocks read by host	B	1) Is there a Bottleneck? 5) More controller RAM? 7) Upgrade improve performance?
Host Write Blocks - Total number of blocks written by host	B	1) Is there a Bottleneck? 5) More controller RAM? 7) Upgrade improve performance?

[0038] Performance metrics can be measured anywhere in the data storage system. For instance, samples are taken at point C for a performance metric concerned with the maximum host wait queue size.

25

[0039] Returning to Fig. 3, at step 56, a performance history based upon the user's storage activity is constructed. The performance history involves those parameters reflecting use of the data storage system, including such parameters as the amount of available space, amount of consumed space, I/O demand, and the amounts of mirror and parity type storage. The performance history reveals how the data storage system is performing over a period of time. The construction of the performance history per step 56 can be conducted in either the data storage system 14 or the host computer station 12.

30

[0040] The performance history is then presented at the host computer station 12 to the user at step 58. Preferably, the performance history is graphically displayed on visual display screen 18 in Fig. 1.

35

[0041] Fig. 5 shows an example screen display 100 for performance history. More particularly, the performance of the data storage system in terms of unused memory capacity (measured in gigabytes), RAID level 1 type storage (measured in gigabytes), and I/O access activity (measured in both Kilobytes per second and I/Os per second) are provided over a three day period. Notice that the performance history shows considerable activity during daytime hours of 6:00 am and 12:00 pm. The displayed time period can be adjusted by the user.

40

[0042] Now suppose the user wants to evaluate one particular time period. For instance, the user might observe an abnormal glitch or trend in the performance history and wish to examine that period more thoroughly. Another example is that the user might wish to avoid periods of atypical usage while analyzing more normal periods. In either case, the user simply selects an evaluation interval by graphically manipulating a interval region, bounded by vertical lines 102 and 104, that is overlaid on the performance history graph (step 60 in Fig. 3). Each boundary line 102, 104 is independently slidable over the time scale provided at the bottom the screen. The depicted evaluation interval is from 9:00 am to 11:00 pm on May 30.

45

[0043] Once an evaluation interval is selected, the next step 62 (Fig. 3) is to determine whether any performance metric is not "in spec" within the selected evaluation interval. This determination is preferably conducted by a performance analysis utility resident in computer station 12. The determination of step 62 can be made in many different ways. One approach is to average the samples for each performance metric measured over time within the specified evaluation interval to derive representative average performance metrics. These averaged values can then be compared to respective specification thresholds to determine if the performance metrics are "in spec". Apart from averaging, other more complex techniques for combining the samples among multiple performance metrics and filtering them to derive composite metrics may be used.

50

[0044] The following table 2 lists two example metric determinations. The "SCSI Queue" result is determined from a simple evaluation of the performance metric "SCSI Queue Max", whereas the "Bandwidth" result is an example composite metric which consists of a combination of raw performance metrics "Host Read Blocks", "Host Write Blocks", and "Active Time" that have been sampled.

TABLE 2

Name	Metric	Spec	Out of Spec	Decision Supported	
5	SCSI Queue	SCSI Queue Max	0-1	≥ 3	4) Additional Disks Help?
	Bandwidth	$\{(Host\ Read\ Blocks + Host\ Write\ Blocks) * 512\} + Active\ Time$	10^6	10^7	1) Is there a Bottleneck?

10 [0045] Another approach is to ascertain an extent to which the performance metrics are "in spec" or "out of spec". This can be accomplished by providing multiple thresholds or tolerance computations for each performance metric which have different degrees of compliance associated therewith. The sampled performance metric is compared to the various thresholds to determine the degree of noncompliance based upon which of the thresholds or tolerance computations the sampled performance metric satisfies. Different conclusions regarding operation of the data storage system might be drawn depending upon how "out of spec" the performance metric is.

15 [0046] If the performance metrics are in compliance with the specification (i.e., the "yes" branch from inquiry step 64 in Fig. 3), the data storage system is operating at a level of optimal performance. The performance analysis utility resident in computer station 12 returns comments stating that all is well for display on the visual display monitor 18 (step 66). This feedback informs the user that no modifications are necessary.

20 [0047] Conversely, if one or more performance metrics are not in compliance with the specification, but instead are "out of spec" (i.e., the "no" branch from inquiry step 64), the data storage system provides suggestions to the user that would help improve performance (step 68). Such suggestions are chosen from a list of possible suggestions based upon which of the one or more performance metrics are "out of spec", and if capable, the degree to which they are "out of spec".

25 [0048] One way to arrive at the appropriate suggestions is to list the predetermined suggestions in a table and correlate them with the performance metric. Table 3 provides an example correlation between the performance metric determinations "SCSI Queue" and "Bandwidth" from Table 2 (as well as others) and the suggestions.

TABLE 3

Out of Spec	In Spec	Suggestions	
30	SCSI Queue	Max Drives	Heavy workload for drives. If condition persists, add drives.
	SCSI Queue Max Drives		Heavy workload for storage system. If condition persists, consider adding more storage systems or upgrading to larger storage system.
35	Bandwidth Host Setting		Host data bandwidth may be inhibiting optimal performance. Adjust host setting.
	Bandwidth	Host Setting	Storage system has inadequate bandwidth for host demand. Upgrade to higher end storage system.

40 [0049] Once one or more suggestions are chosen given the conditions of the performance metrics, the suggestions are displayed to the user via the host computer visual display monitor 18 (step 70). An example performance planning suggestions screen 106 is shown in Fig. 6. This screen relates to the performance history shown in Fig. 5. This screen displays a suggestion "More DRAM would improve performance", and further lists the indications that led to that suggestion. Notice that the cache utilization metric is "out of spec", indicating that a write cache is too full. It is further noted that the DRAM installed is not at maximum. The other indications show normal operation. Based upon these findings, the conclusion of adding more DRAM to improve performance is suggested. From this suggestion, the user can now make an informed decision to add more DRAM or forego adding more DRAM.

45 [0050] If the user would like to evaluate the raw performance metrics which underlie the suggestion to add more DRAM, the next step 72 (Fig. 3) is to display the performance metrics themselves on the visual display monitor 18. An example screen 108 of performance metrics is shown in Fig. 7. Performance metrics concerning "Relocate", "Working Set", "Cache Utilization", "Drive Q", and "Read Multi" are presented. The same evaluation interval from 9:00 am to 11:00 pm on May 30 is bounded by vertical lines 102 and 104.

50 [0051] This bar graph provides the processed data samples extracted during this period of time to provide positive support for the suggestions being provided to the user. Notice, for example, that the cache utilization metric is relatively high in the evaluation interval, and thus yields a finding of "out of spec" which led to the conclusion of adding more DRAM. The user might wish to evaluate various intervals to see if the new averaged metric is still "out of spec".

55 [0052] Notice that each screen in Figs. 5-7 include buttons 110, 112, and 114 to quickly reference the performance

history screen 100, the suggestions screen 106, or the performance metrics screen 108, respectively. These three screens provide the user with a wealth of information to help him or her identify areas for performance improvement in the data storage system. They also assist the user in anticipating future storage requirements for timely upgrades.

[0053] In compliance with the statute, the invention has been described in language more or less specific as to structural and methodical features. It is to be understood, however, that the invention is not limited to the specific features shown and described, since the means herein disclosed comprise preferred forms of putting the invention into effect. The invention is, therefore, claimed in any of its forms or modifications within the proper scope of the appended claims appropriately interpreted in accordance with the doctrine of equivalents.

10

Claims

1. A method for identifying areas for performance improvement in a data storage system (14), the method comprising the following steps:

15

sampling a performance metric during operation of the data storage system;

20

presenting to a user a performance history that is indicative of how the data storage system is performing over a period of time;

25

permitting the user to select an evaluation interval of the performance history;

determining whether the performance metric is at a level for optimal performance of the data storage system for the selected evaluation interval; and

30

based upon said determination of the performance metric for the selected evaluation interval providing at least one suggestion to the user that would help improve performance of the data storage system in the event that the performance metric is not at the level for optimal performance.

35

2. A method according to claim 1 wherein the presenting step comprises graphically displaying the performance history (100) on a display screen (18) and the permitting step comprises permitting the user to select an evaluation interval of the performance history by allowing graphical manipulation of an interval region (102, 104) overlaid on the performance history (100).

40

3. A method according to claim 1 wherein the determining step comprises averaging the samples of the performance metric measured over time to derive a representative average performance metric.

4. A method according to claim 1 wherein the determining step comprises:

45

combining the samples of the performance metric measured over time; and
filtering the combined samples to derive a representative value indicative of the performance metric.

5. A method for identifying areas for performance improvement in a data storage system (14), the method comprising the following steps:

50

defining multiple performance metrics that can be measured during operation of the data storage system (14);

55

providing a specification for the performance metrics whereby the data storage system is operating at optimal performance when the performance metrics are in compliance with the specification;

sampling the performance metrics during operation of the data storage system;

constructing a performance history that is indicative of how the data storage system is performing over a period of time based upon storage demands imposed by a user;

presenting the performance history to the user;

enabling the user to select an evaluation interval of the performance history;;

determining whether any performance metric is not in compliance with the specification for the selected evaluation interval; and

5 in the event that one or more performance metrics are not in compliance with the specification, providing suggestions that would help improve performance of the data storage system based upon which of the one or more performance metrics are not in compliance with the specification for the selected evaluation interval.

6. A method according to claim 5 wherein the specification includes a threshold for individual performance metrics; and the determining step comprises averaging individual performance metrics for the evaluation interval and comparing the averaged performance metrics with their respective thresholds.

10 7. A method according to claim 5 wherein the determining step comprises ascertaining an extent to which one or more performance metrics are not in compliance with the specification.

15 8. A method according to claim 5 wherein the specification includes multiple thresholds for individual performance metrics; and

the determining step for individual performance metrics comprises:

20 averaging the performance metric for the evaluation interval;

comparing the averaged performance metric to the multiple thresholds; and

25 computing degree of noncompliance based upon which of the multiple thresholds the performance metric exceeds.

26 9. A method according to claim 5 wherein:

30 the presenting step comprises graphically displaying the performance history (100) on a display screen (18); and

35 the enabling step comprises graphically overlaying an interval region (102,104) on the performance history and permitting the user to select the evaluation interval by graphically manipulating the interval region.

40 10. A method according to claim 5 further comprising:

35 listing multiple predetermined suggestions in a table which correlates the suggestions to the performance metrics; and

45 the suggestions providing step comprises extracting one or more suggestions from the table based upon which of the one or more performance metrics are not in compliance with the specification.

Patentansprüche

45 1. Ein Verfahren zum Identifizieren von Bereichen zur Leistungsverbesserung in einem Datenspeichersystem (14), wobei das Verfahren die nachfolgenden Schritte aufweist:

50 Abtasten eines Leistungsmaßes während eines Betriebs des Datenspeichersystems;

55 Präsentieren einer Leistungsgeschichte einem Anwender, die anzeigt, wie das Datenspeichersystem über einen Zeitraum hinweg arbeitet;

Ermöglichen, daß der Anwender, ein Bewertungsintervall der Leistungsgeschichte auswählt;

60 Bestimmen, ob das Leistungsmaß auf einer Ebene für optimale Leistung des Datenspeichersystems für das ausgewählte Bewertungsintervall liegt; und

65 Liefert zumindest eines Vorschlags zu dem Anwender basierend auf der Bestimmung des Leistungsmaßes

für das ausgewählte Bewertungsintervall, der helfen würde, die Leistung des Datenspeichersystems zu verbessern, für den Fall, daß das Leistungsmaß nicht auf der Ebene für optimale Leistung liegt.

2. Ein Verfahren gemäß Anspruch 1, bei dem der Schritt des Präsentierens ein graphisches Anzeigen der Leistungsgeschichte (100) auf einem Anzeigebildschirm (18) aufweist und der Schritt des Ermöglichen das Ermöglichen aufweist, daß der Anwender, ein Bewertungsintervall der Leistungsgeschichte auswählt, durch Ermöglichen einer graphischen Manipulation eines Intervallbereichs (102, 104), der der Leistungsgeschichte (100) überlagert ist.
3. Ein Verfahren gemäß Anspruch 1, bei dem der Schritt des Bestimmens ein Mitteln der Abtastwerte des Leistungsmaßes aufweist, die im Laufe der Zeit gemessen wurden, um ein repräsentatives Durchschnittsleistungsmaß abzuleiten.
4. Ein Verfahren gemäß Anspruch 1, bei dem der Schritt des Bestimmens folgende Schritte aufweist.
 - 15 Kombinieren der Abtastwerte des Leistungsmaßes, die im Laufe der Zeit gemessen wurden; undFiltern der kombinierten Abtastwerte, um einen repräsentativen Wert abzuleiten, der das Leistungsmaß angibt.
 5. Ein Verfahren zum Identifizieren von Bereichen zur Leistungsverbesserung in einem Datenspeichersystem (14), wobei das Verfahren folgende Schritte aufweist:
 - 20 Definieren von mehreren Leistungsmaßen, die während des Betriebs des Datenspeichersystems (14) gemessen werden können;
 - 25 Liefert einer Spezifikation für die Leistungsmaße, wobei das Datenspeichersystem mit optimaler Leistung arbeitet, wenn die Leistungsmaße mit der Spezifikation übereinstimmen;
 - 30 Abtasten der Leistungsmaße während des Betriebs des Datenspeichersystems;
 - 35 Erstellen einer Leistungsgeschichte, die angibt, wie das Datenspeichersystem basierend auf Speicheranforderungen, die durch einen Anwender gestellt werden, über einen Zeitraum hinweg arbeitet;
 - 40 Präsentieren einer Leistungsgeschichte dem Anwender;
 - 45 für den Fall, daß ein oder mehrere Leistungsmaße nicht mit der Spezifikation übereinstimmen, Liefert von Vorschlägen, die, basierend darauf, welches der ein oder mehreren Leistungsmaße nicht mit der Spezifikation für das ausgewählte Bewertungsintervall übereinstimmen, helfen würden, die Leistung des Datenspeichersystems zu verbessern.
 6. Ein Verfahren gemäß Anspruch 5, bei dem die Spezifikation einen Schwellenwert für individuelle Leistungsmaße umfaßt; und der Schritt des Bestimmens ein Mitteln individueller Leistungsmaße für das Bewertungsintervall und ein Vergleichen der gemittelten Leistungsmaße mit ihren entsprechenden Schwellenwerten aufweist.
 7. Ein Verfahren gemäß Anspruch 5, bei dem der Schritt des Bestimmens ein Ermitteln eines Ausmaßes aufweist, in dem ein oder mehrere Leistungsmaße nicht mit der Spezifikation übereinstimmen.
 8. Ein Verfahren gemäß Anspruch 5, bei dem die Spezifikation mehrere Schwellenwerte für individuelle Leistungsmaße umfaßt; und der Schritt des Bestimmens für individuelle Leistungsmaße folgende Schritte aufweist:
 - 55 Mitteln des Leistungsmaßes für das Bewertungsintervall;
 - Vergleichen des gemittelten Leistungsmaßes mit den mehreren Schwellenwerten; und

Berechnen des Grads der Nichtübereinstimmung basierend darauf, welchen der mehreren Schwellenwerte das Leistungsmaß übertrifft.

9. Ein Verfahren gemäß Anspruch 5, bei dem:

5

der Schritt des Präsentierens ein graphisches Anzeigen der Leistungsgeschichte (100) auf einer Bildschirm-
anzeige (18) aufweist; und

10

der Schritt des Ermöglichen ein graphisches Überlagern eines Intervallbereichs (102, 104) auf die Leistungs-
geschichte und ein Ermöglichen, daß der Anwender das Bewertungsintervall durch graphisches Manipulieren
des Intervallbereichs auswählt, aufweist.

10. Ein Verfahren gemäß Anspruch 5, das ferner folgende Schritte aufweist:

15

Auflisten mehrerer vorbestimmter Vorschläge in einer Tabelle, die die Vorschläge mit den Leistungsmaßen
korreliert; und

20

der Schritt des Liefern von Vorschlägen ein Extrahieren von einem oder mehr Vorschlägen aus der Tabelle
basierend darauf, welches der ein oder mehreren Leistungsmaße nicht mit der Spezifikation übereinstimmt,
aufweist.

Revendications

25 1. Procédé pour identifier des zones pour l'amélioration des performances dans un système de stockage de données
(14), le procédé comprenant les étapes suivantes :

échantillonner une métrique de performances pendant le fonctionnement du système de stockage de
données ;
30 présenter à un utilisateur un historique des performances indiquant la façon dont le système de stockage de
données fonctionne pendant une certaine période de temps ;
permettre à l'utilisateur de sélectionner un intervalle d'évaluation de l'historique de performances ;
déterminer si la métrique de performances est à un niveau de performances optimales du système de stockage
de données pour l'intervalle d'évaluation sélectionné ; et
35 en se basant sur ladite détermination de la métrique de performances pour l'intervalle d'évaluation sélectionné,
fournir au moins une suggestion à l'utilisateur l'aider à améliorer les performances du système de stockage
de données dans le cas où la métrique de performances n'est pas au niveau des performances optimales.

40 2. Procédé selon la revendication 1, dans lequel l'étape de présentation comprend l'affichage graphique de l'historique
de performances (100) sur un écran d'affichage (18) et l'étape d'autorisation comprend l'autorisation à l'utilisateur
de sélectionner un intervalle d'évaluation de l'historique de performances en permettant une manipulation
graphique d'une région d'intervalle (102, 104) superposée à l'historique de performances (100).

45 3. Procédé selon la revendication 1, dans lequel l'étape de détermination comprend le moyennage des échantillons
de la métrique de performances mesurée dans le temps pour déterminer une métrique de performances moyenne
représentative.

4. Procédé selon la revendication 1, dans lequel l'étape de détermination comprend :

50 la combinaison des échantillons de la métrique de performances mesurée dans le temps ; et
le filtrage des échantillons combinés pour déterminer une valeur représentative indiquant la métrique de performances.

55 5. Procédé pour identifier des zones pour l'amélioration des performances dans un système de stockage de données
(14), le procédé comprenant les étapes suivantes :

définir des métriques de performances multiples pouvant être mesurées pendant le fonctionnement du système de stockage de données ;

5 fournir une spécification pour les métriques de performances telle que le système de stockage de données fonctionne pour des performances optimales lorsque les métriques de performances sont conformes à la spécification ;

10 échantillonner les métriques de performances pendant le fonctionnement du système de stockage de données ; construire un historique des performances indiquant la façon dont le système de stockage de données fonctionne pendant une période de temps basée sur les demandes de stockage imposées par un utilisateur ; présenter l'historique des performances à l'utilisateur ; permettre à l'utilisateur de sélectionner un intervalle d'évaluation de l'historique de performances ; déterminer si une quelconque métrique de performances n'est pas conforme à la spécification pour l'intervalle d'évaluation sélectionné ; et dans le cas où une ou plusieurs métriques de performances ne sont pas conformes à la spécification, fournir des suggestions facilitant l'amélioration des performances du système de stockage de données en se basant sur celle(s) de la ou des métriques de performances qui n'est ou ne sont pas conforme(s) à la spécification pour l'intervalle d'évaluation sélectionné.

15

20 6. Procédé selon la revendication 5, dans lequel la spécification comporte un seuil pour des métriques de performances individuelles ; et l'étape de détermination comprend le moyennage des métriques de performances individuelles pour l'intervalle d'évaluation et la comparaison des métriques de performances moyennées à leurs seuils respectifs.

25 7. Procédé selon la revendication 5, dans lequel l'étape de détermination comprend la vérification de l'importance dont une ou plusieurs métriques de performances ne sont pas conformes à la spécification.

30 8. Procédé selon la revendication 5, dans lequel la spécification comporte des seuils multiples pour des métriques de performances individuelles ; et l'étape de détermination des métriques de performances individuelles comprend :

35 le moyennage de la métrique de performances pour l'intervalle d'évaluation ; la comparaison de la métrique de performances moyennée aux seuils multiples ; et le calcul du degré de non-conformité en se basant sur celui des seuils multiples que dépasse la métrique de performances.



40 9. Procédé selon la revendication 5, dans lequel :

35 l'étape de présentation comprend l'affichage graphique de l'historique de performances (100) sur un écran d'affichage (18) ; et l'étape d'activation comprend la superposition graphique d'une région d'intervalle (102, 104) sur l'historique de performances et la permission à l'utilisateur de sélectionner l'intervalle d'évaluation en manipulant graphiquement la région d'intervalle.


45 10. Procédé selon la revendication 5, comprenant en outre :


45 la fourniture de la liste de multiples suggestions prédéterminées dans une table effectuant la corrélation entre les suggestions et les métriques de performances ; et l'étape de fourniture de suggestions comprend l'extraction d'une ou plusieurs suggestions de la table en se basant sur celle(s) de la ou des métriques de performances qui n'est ou ne sont pas conforme(s) à la spécification.


50



IE II a 2







II II II 5


Performance Planning Suggestions

More DRAM would improve performance

Indications

The drives are balanced on the interval SCSI bus. The relocation metric is in spec. The working set is not changing too rapidly. The working set metric is in spec. There is enough RAID 1 space for the working set. The cache utilization metric is out of spec. The write cache is too full. The drive Q metric is in spec. There are enough drives for the working set. The read multi metric is in spec. There are enough overlapped I/Os. The DRAM installed is not the maximum possible.

Based on interval from 05/30 at 9:00am through 05/30 at 11:00pm



四
五

